

WHAT IS CLAIMED IS:

1. A computer-readable medium having computer-executable instructions, comprising:

transmitting an interval vector from a first member of a  
5 replica set to a second member of the replica set, the  
interval vector comprising one or more intervals of versions,  
each interval having an upper bound and a lower bound.

determining which resources are out-of-sync between the  
members via the interval vector;

10 synchronizing at least one of the out-of-sync resources;  
updating the interval vector to indicate that the at  
least one of the out-of-sync resources is synchronized.

2. The computer-readable medium of claim 1, wherein  
15 determining which resources are out of sync comprises finding  
resources on the second member, each found resource having a  
version not contained in the intervals of the interval vector.

3. The computer-readable medium of claim 2, wherein a  
20 version is not contained in the interval vector if none of the  
intervals in the interval vector include the version.

4. The computer-readable medium of claim 1, wherein each version is a number.

5. The computer-readable medium of claim 1, wherein the intervals of the interval vector are represented as nodes of a binary search tree.

6. The computer-readable medium of claim 5, wherein the binary search tree comprises an AVL tree, a red-black tree, or a splay tree.

7. The computer-readable medium of claim 5, wherein each non-terminal node of the binary search tree has at least one child, the child having an interval with an upper bound less than the lower bound of the interval in the child's parent node or having an interval with a lower bound greater than the upper bound of the child's parent node.

8. The system of claim 5, wherein updating the interval vector to indicate that the at least one of the out-of-sync resources is synchronized comprises reducing the number of nodes in the binary search tree.

9. The computer-readable medium of claim 1, further comprising transmitting another interval vector from the second member of the replica set to the first member of the replica set, the other interval vector serving to distinguish updates from conflicts.

10. The computer-readable medium of claim 1, wherein the out-of-sync resources are organized in a sequence determined by their corresponding version number and wherein synchronizing the at least one of the out-of-sync resources comprises transmitting the resources in an order other than the sequence.

11. The computer-readable medium of claim 1, wherein the interval vector is represented by an interval decision diagram having internal and leaf nodes, wherein internal nodes represent binary digits in a binary expansion of a version sequence number and have two outgoing edges.

20

12. The computer-readable medium of claim 11, wherein the interval decision diagram remains the same size regardless of the number of intervals in the interval vector.

5        13. An apparatus for replicating resources, comprising:  
storage arranged to store the resources, each resource associated with a version;

storage arranged to store an interval vector, the interval vector comprising one or more intervals of versions,  
10 each interval having an upper bound and a lower bound;

a communications mechanism arranged to transmit the interval vector and to receive updates in response thereto, wherein determining whether to update a resource comprises finding that the version associated with the resource is not  
15 included in any of the intervals of the interval vector.

14. The apparatus of claim 13, wherein one of the updates comprises a deletion of a resource.

20        15. The apparatus of claim 13, wherein one of the updates comprises a modification of a resource.

16. The apparatus of claim 13, wherein the interval vector is represented as a binary search tree.

17. The apparatus of claim 13, wherein the interval  
5 vector is represented as a trie embodied as a directed acyclic graph.

18. The apparatus of claim 17, wherein the directed acyclic graph is an interval decision diagram.

10

19. The apparatus of claim 18, wherein the interval decision diagram has a size that remains constant regardless of the number of intervals in the interval vector.

20. A system for replicating resources, comprising:

a first member arranged to participate in a replica set,  
the first member having an interval vector that includes at  
least two intervals, each interval having an upper bound and a  
5 lower bound; and

a second member arranged to participate in the replica  
set,

wherein the first member is arranged to perform acts,  
comprising:

10 transmitting the interval vector to the second  
member and updating the interval vector based on resources  
received from the second member, and

wherein the second member is arranged to perform acts,  
comprising:

15 receiving the first interval vector;

determining which versions of resources the first  
member is missing by finding versions of the resources on the  
second member that are not included in any interval of the  
interval vector; and

20 sending at least one of the resources that has a  
version that is missing to the first member.

21. The system of claim 20, wherein each resource comprises a file.